arXiv:1411.2182v1 [stat.AP] 9 Nov 2014

A Censored Bayesian Hierarchical Model For Precipitation

Yang Liu[†,*], Philip Kokic[†] and K.Shuvo Bakar [‡]
[†]CSIRO, Canberra, Australia.
[‡]Yale University, New Haven, USA.

Version of September 15, 2014 (first revision)

Under review at *Environmetrics*

Abstract

Modelling of precipitation, including extremes, is important for hydrological and agricultural applications. Traditionally, because of large sample properties for data over a large threshold value, generalised Pareto (GP) distributions are often used for modelling extreme rainfall. It can be shown that under certain conditions the generalised hyperbolic (GH) distributions can approximate the power law decay of the GP distribution in the tails. Given their flexible form, this raises the possibility that distributions from the GH family serve as a model for the entire rainfall distribution thus avoiding the need to select a threshold. In this paper, we use a flexible censored hierarchical model that leverages the GH distribution to accommodate data subject to heavy tails and an excessive number of zeros. The fitted model allows estimation of probabilities and return periods of the rainfall extremes, and it produces narrower credible intervals in the tails than the traditional GP method. The model not only fits the tails of the rainfall distribution, but fits the whole distribution very well. It also efficiently represents short-term dependencies in the data so it is suitable for evaluating duration over and below thresholds as well as duration of zero rainfall.

Keywords: Bayesian inference, generalised hyperbolic processes, hierarchical models, rainfall modelling.

1 Introduction

Currently, the General Circulation Model (GCM) is still the most reliable tool for generating the future climate change scenarios (Ye and Li, 2011). But it is widely known that physical models are inadequate for extremes, due to its coarse spatial resolution and the current incomplete understanding of the climate system. Ye and Li (2011) suggest that the poor performance in regional/local precipitation simulation makes it difficult to directly use GCM outputs in climate change impact on extreme precipitation change studies, because extreme precipitation event is most likely a localised phenomenon. For these reasons, statistical models are often considered when modelling rainfall and rainfall extremes.

In fact, statistical modelling of rainfall has important applications in many different fields of

[*]Corresponding Author. Email: Yang.Liu@csiro.au

research including hydrology, agriculture, and environmental sciences. Within each of these fields information is required on a number of spatial and temporal scales. In hydrological applications the frequency and duration of extreme rainfall events over short time periods is very important (Shao et al., 2013; Fowler et al., 2007; Tetzlaff et al., 2005). For agricultural modelling climate information is required on the short-term variations associated with extreme and non-extreme events through the realistic simulation of observed data (Keating et al., 2003; Kokic et al., 2013), as well as intermediate variations associate with seasonal and inter-seasonal variations, and long-term variations due to causes such as climate change. Environmental science applications often require rainfall data with a high degree of spatial resolution (Ashcroft et al., 2011; Ashcroft and Gollan, 2012). For these reasons, statistical approaches that can provide information containing these varying temporal and spatial characteristics is valuable. To address these requirements we need a very flexible statistical approach that can accurately represent upper and lower extremes, skewed data and varying shapes of the rainfall distribution, as well as reliable estimates of the serial dependencies in these data. The objective of this paper is to describe a unified statistical model that moves towards meeting several of these multiple objectives.

This paper illustrates the potential advantages of using a censored generalised hyperbolic (GH) model for rainfall modelling. For extreme rainfall modelling it will be compared to the frequently used approach of fitting a generalised Pareto (GP) distribution to data above a threshold. This approach has been widely used for modelling extremes because the GP distribution is the limiting distribution of statistically independent observations lying above a very large threshold (Smith, 1987). GP distributions have been extensively used as a primary tool for modelling extreme values in a wide range of disciplines. Among others, Van Montfort and Witter (1986), Coles et al. (2003), Li et al. (2005) and Yiou et al. (2008) use the technique to model the tails of rainfall and temperature distributions, Gilli et al. (2006) use the GP distribution to model extreme financial returns over a large threshold, and Finkenstadt and Rootzén (2004) discuss the applications of similar methods in insurance, telecommunications and environmental science. The rationale behind considering a GH model is that the GH distribution has an extremely flexible form, so that it can accommodate datasets with heavy tails and skewness. In this paper it is shown that GH distributions can approximate the power law decay of the GP distribution in the tails, as well as achieve a good fit at the centre and shoulders of the distribution. Fitting a GH model does not involve choosing an appropriate threshold which can be a major drawback of GP models because they do not account for the uncertainty associated with the choice of threshold, and potentially information below the threshold is discarded for estimating extremes. This is a significant advantage of modelling with the GH distribution as it allows us to make inferences of not only the rainfall extremes, but also the entire time series. It is also advantageous because the credible intervals (CI) for return periods of rainfall extremes obtained under the GH model have a larger sample contributing to the estimate.

We also compare the GH distribution to the Generalised Weibull (GW) distribution (Mudholkar and Huston, 1996). This distribution, also known as the extended Burr XII distribution, has been used for modelling extreme values of rainfall (Shao et al., 2004, 2013). It has also been widely applied to areas such as economics and insurance (Maddala, 1986; Kleiber and Kotz, 2003). In this paper, it is shown that due to the more flexible shape of the GH disribution, the GH model has the potential to outperform the GW distribution in its ability to fit the entire rainfall distribution.

The underlying model we develop for rainfall is an ARMA process with GARCH errors, whose

innovations are independent and identically distributed GH random variables. The ARMA and GARCH processes account for the autocorrelation and heteroscedasticity of the time series, whereas the GH distribution addresses the skewed and heavy-tailed nature of rainfall data. Zero rainfall contained in the time series is considered as censored observations from the underlying distribution; an approach that has been employed by various authors to represent dry days in stochastic rainfall models (Hutchinson, 1995; Wang et al., 2012; Sahu et al., 2010). We apply the Metropolis-Hastings algorithm (Metropolis et al., 1953) for sampling the joint posterior of model parameters. This enables the easy production of CIs of return periods and other important statistics using Bayesian Markov Chain Monte Carlo (MCMC) sampling.

This paper also accounts for various climatic drivers that have influence on precipitation levels. Understanding the effects of climatic drives on rainfall and the pattern of seasonal precipitation is important for agricultural and water risk-management. Empirical evidences on the relationships between them are discussed in many studies (Bakar and Sahu, 2013; Bakar and Kokic, 2014; K.S.Bakar and Jin, 2014; Kokic et al., 2013). Three important climatic drivers are the El Nino southern oscillation anomaly (NINO 3.4), Southern Hemisphere Annular Mode (SAMI) and Indian Ocean Dipole (IOD). It has been observed that these indices have potential teleconnection with precipitation levels.

This paper is divided into several sections. In the next section, a brief review of the GH, GP and GW distribution is presented. In Section 3, we describe the general setup of the ARMA-GARCH-GH model and model assumptions. In Section 4, we illustrate how to use the censored ARMA-GARCH-GH model for modelling rainfall data. In Section 5 and 6, a simulation study of the ARMA-GARCH-GH model and a sensitivity analysis are presented. In Section 7 we apply the proposed model to rainfall data from three locations (Brisbane, Pardelup and the Oaks), which are chosen in order to test the model in different climatic zones. Finally it will be compared with the GP and GW models. Section 9 includes a short discussion and future directions.

2 background

2.1 The Generalised Pareto and Other Related Distributions

The generalised Pareto (GP) distribution, whose density is given by

$$d_{GP(\kappa,\mu,\sigma)}(x) = \frac{1}{\sigma}\left(1+\frac{\kappa(x-\mu)}{\sigma}\right)^{-\left(\frac{1}{\kappa}+1\right)} \tag{1}$$

is often used for extreme value modelling. This idea was first suggested by Pickands III (1971) and has been developed by, among others, DuMouchel (1983), Smith (1984), Smith (1987), Hosking and Wallis (1987) and Joe (1987). In particular, Davison and Smith (1990) used a generalised Pareto distribution as the modelling distribution for exceedances over high thresholds. More recently, Li et al. (2005) applied this approach, coupled with a threshold value selection criterion suggested by Coles et al. (2001), to model extreme rainfall.

Another relevant distribution is the Burr distribution. Since its introduction by Burr (1942), the Burr distribution has been studied intensively and widely used in many disciplines, particularly economics and actuarial science (Maddala (1986), Kleiber and Kotz (2003) and Champernowne

(1952)). The density function of the Burr distribution is given by

$$f_{Burr(c,k,\sigma)}(x) = ck\frac{(x/\sigma)^{c-1}}{(1+(x/\sigma)^c)^{k+1}},\ x>0,\ \text{for}\, c,k,\sigma>0. \tag{2}$$

We notice that the Burr distribution is in the Pareto family and is a special case of the Pareto type-IV. Thus the Burr distribution is capable of capturing skewness and heavy tails. To further improve its flexibility, the generalised Weibull (GW) distribution (also known as the extended Burr distribution) was introduced by Mudholkar and Huston (1996). The density is given by

$$f_{GW(c,r,\nu)}(x) = \begin{cases} \frac{c}{\nu}\left(\frac{x}{\nu}\right)^{c-1}\left(1-r\left(\frac{x}{\nu}\right)^c\right)^{\frac{1}{r-1}},\ r\neq 0, \\ \frac{c}{\nu}\left(\frac{x}{\nu}\right)^{c-1}\exp\left(-\left(\frac{x}{\nu}\right)^c\right),\ r=0. \end{cases} \tag{3}$$

We notice that for $r\neq 0$, $GW(c,r,\nu)$ is just the Burr distribution with the reparameterisation $r=-1/k$ and $\nu=\sigma/k^{1/c}$; for $r=0$, (3) reduces to the Weibull distribution. The GW distribution was also used by Shao et al. (2004) and Shao et al. (2013) for extreme flood modelling and dynamic hydrological modelling.

2.2 The Generalised Hyperbolic Distribution

The generalised hyperbolic (GH) distribution was first introduced by Barndorff-Nielsen (1977) in connection with dune movements modelling. Their Lebesgue density is defined as

$$\begin{aligned} d_{GH(\lambda,\chi,\psi,\mu,\Sigma,\gamma)}(x) =& \frac{\left(\sqrt{\frac{\psi}{\chi}}\right)^{\lambda}\left(\psi+\gamma^2\Sigma^{-1}\right)^{\frac{1}{2}-\lambda}}{\sqrt{2\pi|\Sigma|}K_\lambda\left(\sqrt{\chi\psi}\right)} \\ & \frac{K_{\lambda-\frac{1}{2}}\left(\sqrt{\left(\chi+(x-\mu)^2\Sigma^{-1}\right)\left(\psi+\gamma^2\Sigma^{-1}\right)}\right)\exp\left((x-\mu)\Sigma^{-1}\gamma\right)}{\left(\sqrt{\left(\chi+(x-\mu)^2\Sigma^{-1}\right)\left(\psi+\gamma^2\Sigma^{-1}\right)}\right)^{\frac{1}{2}-\lambda}}, \end{aligned} \tag{4}$$

where K_ν is the modified Bessel function of the third kind with index ν, $\lambda\in\mathbb{R}$, $\chi,\psi\in\mathbb{R}^+$ and $\mu,\Sigma,\gamma\in\mathbb{R}$. To gain some intuition on the above expression, one often writes the generalised hyperbolic distribution as the following mean-variance mixture. A random variable X is said to have a GH distribution if

$$X=\mu+\gamma W+\sqrt{\Sigma}\sqrt{W}Z, \tag{5}$$

where Z is a Normal random variable with zero mean and unit variance, and W has a generalised inverse Gaussian distribution with parameters λ,χ and ψ, or $W\sim GIG(\lambda,\chi,\psi)$, whose density function is given by (41). Since X given $W=w$ is Normal with conditional mean $\mu+\gamma w$ and variance Σw, it is clear that μ and Σ are location and dispersion parameters, respectively. There is a further scale parameter χ, a skewness parameter ψ to allow for flexible tail modelling; and the scalar λ, which characterises certain subclass and also influences the size of mass contained in the

tails.

One of the appealing properties of normal mixtures is that the moment generating function of a GH random variable X can be easily calculated using the moment generating function of the generalised inverse Gaussian distribution. In particular, the mean and variance are given by

$$E(X) = \mu + \gamma E(W) \text{ and } Var(X) = \gamma^2 Var(W) + \Sigma E(W). \tag{6}$$

Another attractive property of the GH distribution is that, as the name suggests, it is of a very general form, and contains as special cases many of important distribution widely used in the literature. It includes, among others, the Student's t-distribution, the Laplace distribution, the hyperbolic distribution, the normal-inverse Gaussian distribution and the variance-gamma distribution. It is often used in economics, with particular application in the fields of modelling financial markets and risk management, due to its semi-heavy tails.

Shao et al. (2004) has shown that the GW distribution can approximate the GP distribution in the tails. We now present a similar result for the GH distribution.

Proposition 1. *For any fixed threshold μ, the power law decay of the GP distribution, whose density function is given by* (1) *can be approximated by the GH distribution in the tails if the shape parameter κ of the GP distribution satisfies either of the two conditions:*

1. $\kappa > 0$ *(heavy tail);*
2. $\kappa \to 0$ *(light tail).*

The proof is given in Appendix A.1. We note that the tail for a GP distribution cannot be approximated by the GH distribution if the GP distribution has a bounded tail, i.e. $\kappa < 0$. However this scenario is unrealistic for precipitation data. Therefore Proposition 1 suggests the possibility that distributions from the GH family serve as the underlying distribution for rainfall data.

There are several parameterisations of the GH distribution. The $(\lambda, \chi, \psi, \mu, \Sigma, \gamma)$ parameterisation used in this paper has a drawback of an identification problem (Barndorff-Nielsen and Shephard, 2001), i.e. $GH(\lambda, \chi, \psi, \mu, \Sigma, \gamma)$ and $GH(\lambda, \chi/k, k\psi, \mu, k\Sigma, k\gamma)$ are identical for any $k > 0$. This is because that χ and Σ are not separately identified. Therefore, an identification problem can be an issue when we try to fit the GH distribution to data. This problem can be solved by introducing a suitable constraint on the parameters. Barndorff-Nielsen and Shephard (2001) fixed the dispersion parameter Σ to be 1, or in the multivariate case, the determinant of Σ to be 1. However under their setup, it is difficult to reparameterise the GH distribution so that it has mean zero and unit variance because this involves reparameterising scalars inside convolutions of Bessel functions with different indices. Such standardisation is necessary when we develop the ARMA-GARCH-GH model in the next section. Fortunately, Mencia and Sentana (2004) suggested that the identification problem can be also solved by fixing $\chi = 1$. We will use their result here, because it removes the need for reparameterising the parameter χ inside the Bessel function, hence leads to a simple standardisation (see Appendix A.2 for proof).

Proposition 2. *Let us modify* (4) *by replacing $\gamma = \tau\Sigma$, $\tau > 0$, we have $X \sim GH(\lambda, \chi, \psi, \mu, \Sigma, \tau)$.*

For any λ, ψ, τ *and fixed* $\chi = 1$*, if* μ *and* Σ *satisfy*

$$\mu = -\frac{\tau\Sigma M_{\lambda+1}\left(\sqrt{\psi}\right)}{\sqrt{\psi}} \text{ and } \Sigma = \frac{\sqrt{1+4\left(N_{\lambda+2}\left(\sqrt{\psi}\right)-1\right)\tau^2}-1}{2\tau^2 M_{\lambda+1}\left(\sqrt{\psi}\right)\left(N_{\lambda+2}\left(\sqrt{\psi}\right)-1\right)/\sqrt{\psi}}, \tag{7}$$

where

$$M_{\lambda+1}\left(\sqrt{\psi}\right) = \frac{K_{\lambda+1}\left(\sqrt{\psi}\right)}{K_{\lambda}\left(\sqrt{\psi}\right)} \text{ and } N_{\lambda+2}\left(\sqrt{\psi}\right) = \frac{K_{\lambda+2}\left(\sqrt{\psi}\right)K_{\lambda}\left(\sqrt{\psi}\right)}{K_{\lambda+1}\left(\sqrt{\psi}\right)^2},$$

then $E(X) = 0$ *and* $Var(X) = 1$.

We note that without modifying (4) as in Proposition 2, one could simply write μ and Σ as expressions of λ, ψ and γ, but this does not guarantee the positivity of Σ. This problem can be solved by introducing the parameter τ to create a quadratic function of Σ. After standardisation the original $(\lambda, \chi, \psi, \mu, \Sigma, \gamma)$ parameterisation is used to make use of the **R** package `ghyp` (Breymann and Lüthi, 2013) for simulation.

Unlike models based on the extreme value theory, such as the GP distribution, which disregards data below a large threshold and does not account for the uncertainty associated with the choice of threshold, fitting a GH model does not involve choosing such a threshold. This is a significant advantage of modelling with the GH distribution as it allows us to make inferences of not only the rainfall extremes, but also the entire time series. Consequently the prediction credible intervals (CI) for return periods of rainfall extremes and other statistical inference obtained under the GH model have a larger sample contributing to the estimates. We will now construct our Bayesian hierarchical model and demonstrate these advantages.

3 ARMA-GARCH-GH Model

In this section, a formulation of the proposed statistical model for the underlying distribution of rainfall is presented. Before we introduce the model, there are four important features of rainfall data that need to be addressed, namely autocorrelation, heteroscedasticity, heavy-tailedness and an excessive amount of zero observations. We will discuss the first three here and deal with the zero problem in Section 4.

Let us consider historical rainfall data from the Brisbane climate observation station in south-east Queensland obtained from the Commonwealth Bureau of Meteorology. The sample period is from 1 January 1889, to 22 January 2014, for a total of 45,677 observations, including 22,677 wet season (November–April) observations. The time series of the wet season daily rainfall is shown in Figure 1. The middle panel shows a plot of the autocorrelation function, and indicates that there is autocorrelation present in the data. The right panel shows rainfall volatility during the period of record, from which strong heteroscedasticity can be identified. This may be a consequence of seasonal variation and alternating dry-wet changes. The heavy-tailed nature of the rainfall distribution is illustrated in both the left and right panels, where the majority of observations are scattered below 75 mm, many others distributed between 100 and 150 mm, and a few days with heavy rainfall over 200 mm are observed, which are considered rare or extreme events and are of great interest.

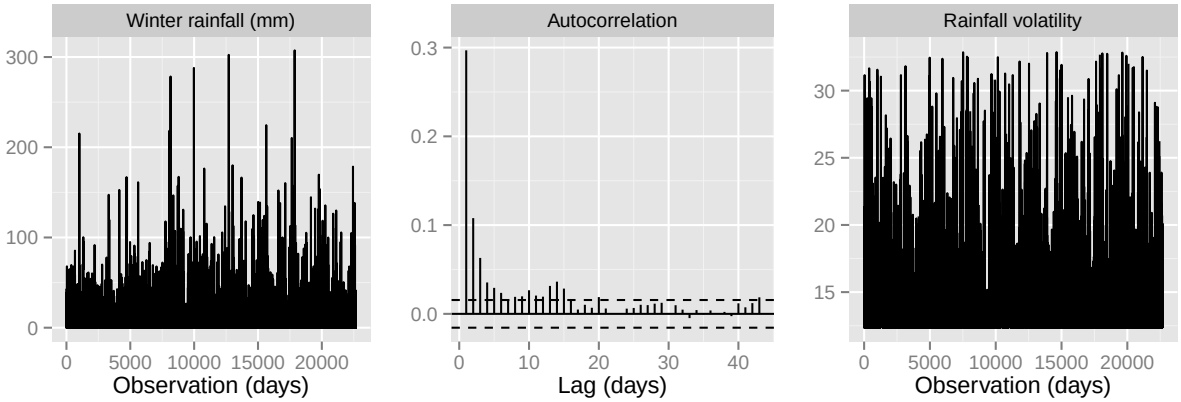

Figure 1: Rainfall data from 1 January 1889, to 22 January 2014 at the weather station Brisbane, Queensland. The left panel is the time series of summer daily rainfall (mm); the middle panel shows a plot of the autocorrelation function, and indicates that there is autocorrelation present in the data; the right panel shows rainfall volatility in the period of record, from which strong heteroscedasticity can be identified. Thus there are sub-samples that have different variabilities from others.

To address the autocorrelation, heteroscedasticity and heavy-tailed features of rainfall data, we propose the following model. An autoregressive and moving average (ARMA, Whittle (1951)) model is used to remove persistence. It generates rainfall residuals with no correlation with past values. A generalised autoregressive conditional heteroscedasticity (GARCH, Bollerslev (1986)) model is also used to account for empirical features in the volatility of rainfall observations. Finally the skewness and heavy-tailed nature of the rainfall distribution can be modelled with a standardised GH distribution. We are now in a position to introduce the ARMA-GARCH-GH model:

$$y_t = a_0 + \sum_{i=1}^{p} a_i y_{t-i} + \sum_{i=1}^{q} b_i \varepsilon_{t-i} + \sum_{i=1}^{M} \beta_i u_{it} + \varepsilon_t \tag{8}$$

$$\varepsilon_t = z_t \sigma_t, \tag{9}$$

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{m} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{r} \omega_i \sigma_{t-i}^2, \tag{10}$$

where y_t, $t = 1,\ldots,T$ is the observed data, u_i is the ith covariate with coefficient β_i, and $z_t \sim GH(\lambda,\chi,\psi,\mu,\Sigma,\gamma)$ are independent and identically distributed standardised GH random variables with zero mean and unit variance. We let $\Phi(L) = \sum_{i=1}^{p} a_i L^i$, $\Theta(L) = \sum_{i=1}^{q} b_i L^i$, $A(L) = \sum_{i=1}^{m} \alpha_i L^i$ and $B(L) = \sum_{i=1}^{r} \omega_j L^j$, where L is the lag operator.

The mean process is given by (8) with autoregressive coefficients a_i for $i = 1,\ldots,p$ and moving average coefficients b_i for $i = 1,\ldots,q$. In a standard ARMA model, however, the conditional variance given the past is constant, so it cannot take into account for heteroscedasticity of the time series. In this context we use the GARCH model given by (9) and (10) in which the conditional variance of innovations ε_t is non constant and depends on time, and therefore can model the randomly varying volatility.

In practice, we impose several constraints on the parameters in the ARMA-GARCH-GH model:

1. All roots of $1 - \Phi(L) = 0$ and $1 + \Theta(L) = 0$ are outside the unit circle;

2. $\alpha_0 > 0$, $\alpha_i > 0$ for $i = 1, \ldots, m$ and $\omega_i > 0$ for $i = 1, \ldots, r$;

3. permitted parameters of GH distributions are $\lambda \in \mathbb{R}$, $\chi, \psi \in \mathbb{R}^+$ and $\mu, \Sigma, \gamma \in \mathbb{R}$. Using Proposition 2 the GH distribution can be standardised so that it has zero mean and unit variance, so Σ and μ can be written as functions of λ, ψ and γ. This requires $\chi = 1$.

Condition 1 is to ensure the stationarity and invertibility of the $ARMA(p,q)$ process; 2 is imposed to guarantee that the conditional variance σ_t^2 is always positive. Bollerslev (1986) shows that the $GARCH(m,r)$ process is weak-sense stationary if and only if $A(1) + B(1) < 1$, and it has mean and variance $E(\varepsilon_t) = 0$ and $Var(\varepsilon_t) = \alpha_0/(1 - A(1) - B(1))$. However, we will not impose the stationarity constraint on the GARCH model since the Bayesian analysis enables us to test variance stationarity condition and estimate the density of the unconditional variance $Var(\varepsilon_t)$ when the condition is satisfied.

4 Censored Bayesian Hierarchical Model For Precipitation

In this section, we describe an application of the above argument and illustrate how to apply the ARMA-GARCH-GH model incorporated with covariates to rainfall data. The rainfall data shows strong autocorrelation and heteroscedasticity that must be adjusted for. Another important feature of the dataset is that it contains an excessive number of zeros, which are considered as censored observations from the underlying process given by (8), (9) and (10).

We begin by describing weekly rainfall sums $X_t = (x_1, \ldots, x_T)$ by the following censored process. Here we only consider $ARMA(1,0)$ and $GARCH(0,1)$, higher order models follow by reproducing carefully the result, keeping track of the time dependence:

$$x_t = y_t I_{\{y_t > 0\}} \tag{11}$$

$$y_t = a_0 + a_1 y_{t-1} + \sum_{i=1}^{M} \beta_i u_{it} + \varepsilon_t \tag{12}$$

$$\varepsilon_t = z_t \sigma_t \tag{13}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2, \tag{14}$$

where u_{it} is the ith regressor with coefficient β_i, $I_{\{\cdot\}}$ is the indicator function, and $z_t \sim GH(\lambda, \chi, \psi, \mu, \Sigma, \gamma), \gamma = \tau\Sigma$ are independent and identically distributed with zero mean and unit variance.

Let us first focus on the uncensored model. Let $\theta = (\lambda, \chi, \psi, \mu, \Sigma, \tau)$, $\varphi = (a_0, a_1, \beta_1, \ldots, \beta_M, \alpha_0, \alpha_1, \theta)$ and $\varphi^* = (\varphi, y_0, \varepsilon_0, \sigma_0)$, where $\sigma_0 = \sqrt{\alpha_0}$, $y_0 = \varepsilon_0$ and ε_0 is the pre sample error. We also assume $\varepsilon_i = \sigma_i = y_i = 0$ for $i < 0$. To perform Bayesian analysis for the ARMA-GARCH-GH model, we construct the posterior density function of the model:

$$\pi(\varphi^* | Y) \propto f(Y | \varphi^*) \pi(\varphi^*), \tag{15}$$

where $f(Y|\varphi^*)$ is the likelihood function and $\pi(\varphi^*)$ is the prior. The likelihood function has the following representation:

$$\begin{aligned} f(Y|\varphi^*) =& f(y_1|\varphi^*,\ldots,y_{-1},y_0)\, f(y_2|\varphi^*,\ldots,y_{-1},y_0,y_1)\, f(y_3|\varphi^*,\ldots,y_{-1},y_0,y_1,y_2),\ldots, \\ & f(y_T|\varphi^*,\ldots,y_{-1},y_0,y_1,\ldots,y_{T-1}). \end{aligned} \tag{16}$$

For any term in (16), say $f(y_t|\varphi^*\ldots,y_{-1},y_0,y_1,\ldots,y_{t-1})$, we know

$$\frac{y_t-\mu_t}{\sigma_t} = z_t \sim GH(\lambda,\chi,\psi,\mu,\Sigma,\gamma),\ \gamma=\beta\Sigma, \tag{17}$$

then $y_t \sim GH\left(\lambda,\chi,\psi,\mu\sigma_t+\mu_t,\Sigma\sigma_t^2,\gamma\sigma_t\right)$, where

$$\mu_t = a_0 + a_1 y_{t-1} + \sum_{i=1}^{M}\beta_i u_{it} \text{ and } \sigma_t = \sqrt{\alpha_0+\alpha_1\varepsilon_{t-1}^2}.$$

If we write $\widetilde{\Sigma}_t = \Sigma\sigma_t^2$, $\widetilde{\gamma}_t = \gamma\sigma_t$ and $\widetilde{\mu}_t = \mu\sigma_t+\mu_t$, the joint posterior is given by

$$\begin{aligned} f(\varphi^*|Y) =& \prod_{t=1}^{T} \frac{\left(\sqrt{\frac{\psi}{\chi}}\right)^{\lambda}\left(\psi+\widetilde{\gamma}_t^2\widetilde{\Sigma}_t^{-1}\right)^{\frac{1}{2}-\lambda}}{\sqrt{2\pi\left|\widetilde{\Sigma}_t\right|}K_\lambda\left(\sqrt{\chi\psi}\right)}\times \\ & \prod_{t=1}^{T} \frac{K_{\lambda-\frac{1}{2}}\left(\sqrt{\left(\chi+(x-\widetilde{\mu}_t)^2\right)\left(\psi+\widetilde{\gamma}_t^2\widetilde{\Sigma}_t^{-1}\right)}\right)\exp\left((x-\widetilde{\mu}_t)\widetilde{\Sigma}_t^{-1}\widetilde{\gamma}_t\right)}{\left(\sqrt{\left(\chi+(x-\widetilde{\mu}_t)^2\right)\left(\psi+\widetilde{\gamma}_t^2\widetilde{\Sigma}_t^{-1}\right)}\right)^{\frac{1}{2}-\lambda}}\times \\ & \pi(a_0)\,\pi(a_1)\prod_{i=1}^{M}\pi(\beta_i)\,\pi(\alpha_0)\,\pi(\alpha_1)\,\pi(\theta). \end{aligned} \tag{18}$$

We use normal priors on the parameters of the mean process, truncated normal priors on GARCH parameters, and normal and truncated normal priors on GH parameters.

Due to the recursive nature of ARMA and GARCH models, the joint posterior and full conditional densities cannot be expressed in closed form. Therefore we cannot use the Gibbs sampler and need to rely on a more elaborated Metropolis-Hasting (MH) simulation strategy to approximate the posterior density (Metropolis et al., 1953; Hastings, 1970). The sampling strategy relies on the construction of a Markov chain with realisations $\varphi^{[0]},\varphi^{[1]},\ldots,\varphi^{[j]},\ldots$ in the parameter space. Under appropriate regularity conditions, asymptotic results guarantee that as $j\to\infty$, $\varphi^{[j]}$ converges in distribution to a random variable whose density is $\pi(\varphi|Y)$. Hence after discarding a burn-in of the first draws, the realised values of the chain can be used to make inference about the joint posterior. In this paper, a modified Metropolis-Hastings algorithm was used for the sampling the joint posterior. We use normal proposals and whose variances are tuned to give an acceptance rate around 30%. The block sampling techique is also used to address the correlation between the GH parameters.

The above algorithm only applies when all y_t are observed, that is, $x_t = y_t$ for all t. Due to the zero-inflated nature of the dataset, however, we are not able to observe the underlying process Y_t directly for all t. Instead we observe Y_t only when it is positive, so X_t is the value we observe due to

censoring. We use the conditional density of y_t given by (18) to simulate censored observations and recover the underlying process.

Suppose $y_1,\ldots,y_{t-1}$ are observed and y_t is censored, i.e. $x_i = y_i$ for $i = 1,\ldots,t-1$, and $x_t = 0$, and let $\varphi^{[j]} = \left(a_0^{[j]}, a_1^{[j]}, \beta_1^{[j]}, \ldots, \beta_M^{[j]}, \alpha_0^{[j]}, \alpha_1^{[j]}, \lambda^{[j]}, \psi^{[j]}, \tau^{[j]}\right)$ be the j^{th} realisation of the Markov chain in the parameter space. To recover the uncensored process Y_t at time t, we sample from the distribution of Y_t conditional on $Y_t < 0$ and given the values $y_{t-1}, y_{t-2}\ldots$, where

$$
\begin{aligned}
Y_t &\sim GH_{(-\infty,0]}\left(\lambda^{[j]}, 1, \psi^{[j]}, \mu^{[j]}\sigma_t + \mu_t, \Sigma^{[j]}\sigma_t^2, \gamma^{[j]}\sigma_t\right), \qquad (19)\\
\Sigma^{[j]} &= \frac{\sqrt{1+4\left(N_{\lambda^{[j]}+1}\left(\sqrt{\psi^{[j]}}\right)-1\right)\left(\tau^{[j]}\right)^2}-1}{2\left(\tau^{[j]}\right)^2 M_{\lambda^{[j]}+1}\left(\sqrt{\psi^{[j]}}\right)\left(N_{\lambda^{[j]}+1}\left(\sqrt{\psi^{[j]}}\right)-1\right)/\sqrt{\psi^{[j]}}},\\
\mu^{[j]} &= -\frac{\tau^{[j]}\Sigma^{[j]} M_{\lambda^{[j]}+1}\left(\sqrt{\psi^{[j]}}\right)}{\sqrt{\psi^{[j]}}},\\
\gamma^{[j]} &= \tau^{[j]}\Sigma^{[j]},\\
\sigma_t &= \sqrt{\alpha_0^{[j]} + \alpha_1^{[j]}\varepsilon_{t-1}^2},\\
\mu_t &= a_0^{[j]} + a_1^{[j]} x_{t-1} + \sum_{i=1}^{M} \beta_i^{[j]} u_{it},
\end{aligned}
$$

and $GH_{(-\infty,0]}$ is the normalised truncated GH distribution to $(-\infty,0]$. For later censored observations x_{t_k}, $t_k > t$, the simulated Y_t is used as an observed value. The data augmentation process is repeated for each MCMC iteration. Therefore, by simulation we are able to recover the underlying process y_t for all t, and for all sets of parameter realisations $\varphi^{[j]}$ in the MCMC chain. This approach to recovering the latent process is quite standard in models involving censoring, see e.g. Sahu et al. (2010).

We produce CIs for return periods by independently simulating time series a large number of times using the MH algorithm at each MCMC iteration and (19) from the ARMA-GARCH-GH model, and then by selecting the 2.5% and 97.5% quantiles from the estimated return periods.

5 Simulation Study

In this section we present a short simulation study for the proposed ARMA-GARCH-GH model. The model is applied to simulated data and the parameter estimates are compared to their true values. The simulated data is comparable to the real rainfall data and it contains 3,200 observations with a censoring rate of around 8%. The simulation result is shown in Table 1. It is evident that the 95% CIs for the parameters contain their true values.

6 Sensitivity Analysis

One of the disadvantages of Bayesian analysis is due to the fact that the prior distribution for the parameters can have a significant impact on the posterior distribution, and consequently, leads to biased results. We check the sensitivity of the model by using diffierent hyper-parameters of the

Parameters	True values	Posterior mean	95% CIs
a_0	5	4.91	4.82, 5.09
a_1	0.5	0.51	0.48, 0.54
α_0	13	13.36	12.50, 14.89
α_1	0.2	0.19	0.16, 0.23
λ	-0.2	-0.17	-0.60, 0.02
τ	15	15.71	9.30, 25.18
ψ	0.25	0.27	0.16, 0.42

Table 1: A simulation study of the ARMA-GARCH-GH model.

Normal and truncated Normal priors. In the case where the original priors are uniform distributions, Beta distributions on the same support are used as alternatives. The results show that our model is not very sensitive to the choice of the hyper-parameter values. For brevity these results have been omitted from the paper.

7 Modelling Results

We apply the ARMA-GARCH-GH model to rainfall data from three locations: Pardelup in south-west Western Australia, The Oaks in south-east New South Wales and Brisbane in south-east Queensland. These locations were chosen in order to test the model in different climatic situations. The first location is in a temperate winter-rainfall dominated zone with a greater number of extreme rainfall events having been observed in recent decades. The second location is on the SW edge of the Sydney basin, and being located close to the coast it experiences a relatively wet temperate climate. However, the area's weather is highly variable: drought and bushfire on the one hand, and storms and heavy rainfall on the other. The last location is a metropolis with a population of more than two million, and in a summer rainfall dominated sub-tropical climate zone. During summer, thunderstorms are common over Brisbane, with the more severe events accompanied by large damaging torrential rain and destructive winds. The city also lies in the tropical cyclone risk area. The last to affect Brisbane was Tropical Cyclone Hamish in March 2009.

Three important climatic drivers the El Nino southern oscillation anomaly (NINO 3.4), Southern Hemisphere Annular Mode (SAMI) and Indian Ocean Dipole (IOD) are considered in this study to understand their effects on precipitation levels. The NINO 3.4 index is a monthly time series of mean sea surface temperature (SST) index from the El Nino region that covers 5 south to 5 north and 125 west to 175 east. The SAMI index includes the Antarctic oscillation and obtained by the differences in the normalised monthly zonal mean seal level pressure between 40 and 70 south. The IOD index used in this study is calculated from monthly SST anomaly, which is a coupled ocean and atmosphere phenomenon in the equatorial Indian Ocean. In addition, for weather station Brisbane we also consider the Norfolk-Hawaii index (NHI). The NHI is an index of the Inter-decadal Pacific Oscillation (IPO). The IPO is a slow background change in Pacific Ocean SSTs, which affects the relationship between the El Nino-Southern Oscillation and Queensland summer rainfall.

Parameters	The Oaks		Brisbane		Pardelup	
	mean	95% CIs	mean	95% CIs	mean	95% CIs
a_0	11.42	10.55,12.40	30.20	29.59,30.89	20.62	18.30,22.65
a_1	0.07	0.01,0.13	0.17	0.01,0.35	0.04	-0.04, 0.14
α_0	152.26	115.38,204.64	2508.63	2296.72,2914.88	271.24	209.20,427.56
α_1	0.47	0.2558,0.77	0.67	0.05,2.27	0.11	0.01,0.28
λ	0.27	0.04,0.59	0.48	0.39,0.57	-1.86	-3.08,-0.59
β	0.49	0.33,0.68	1.36	1.15,1.61	0.51	0.24,0.79
ψ	0.20	0.01,0.49	0.07×10^{-2}	0.00, 0.28×10^{-2}	0.82	0.01,2.8
$\beta 1$ (NINO3.4)	-0.08	-0.24,0.00	-0.45	-1.11,0.41	-0.76	-3.08,1.07
$\beta 2$ (SAMI)	-0.50	-0.92,-0.09	-	-	-0.37	-0.96,0.21
$\beta 3$ (IOD)	-0.03	-0.88,0.84	-	-	-0.74	-1.88,0.39
$\beta 4$ (IPO)	-	-	0.96	0.09,2.33	-	-

Table 2: Posterior mean and corresponding 95% CIs of the parameters of the ARMA-GARCH-GH model.

7.1 Model Based Analysis

We present the results for three locations below and discuss in detail the result for Brisbane, whereas results for Pardelup and the Oaks is presented in Appendix B. The models are fitted with 10,000 MCMC iterations, and the first 3,000 samples are discarded as burn-in. The MCMC chains converged quickly after a few hundred iterations. For brevity, these results and other MCMC diagnostics are omitted.

Table 2 shows the model parameter estimates and their 95% CIs. The autoregressive parameters show a significant positive autocorrelation between successive weeks for rainfall at the Oaks and Brisbane. The GARCH parameter α_1 is significant at all locations, indicating that there are a considerable amount of clustering in weekly rainfall data. As expected, the generalised hyperbolic shape and skewness parameters show that the estimated distribution is highly skewed to the right. A negative and significant coefficient for NINO3.4 and SAMI at the Oaks is consistant with their effect as reported by the Australian Bureau of Meteorology (http://www.bom.gov.au). The sign of the IPO coefficient is consistent with the asymmetric response identified by King et al. (2013).

We should also recognise that the lagged effects of climatic drivers on precipitations can be important (Drosdowsky and Chambers, 2001; Stone et al., 1996). Schepen et al. (2012) argued that concurrent relationships between climatic drivers and precipitations does not imply a lagged relationship, so it is important to understand the effects of lagged climate indices on rainfall. However it is outside our scope, and these exercises have been omitted from the paper.

We can test the variance stationarity condition and estimate the unconditional variance of ε_t when the condition is met. As shown by Bollerslev (1986), under the $GARCH(1,1)$ specification, the process is stationary if $\alpha_1+\beta_1<1$. With a value close to one, past volatility will have a longer impact on the future conditional variance. Under $GARCH(0,1)$, the posterior density of α_1 has a mean of 0.47, 0.57 and 0.11 at three locations. A simple t-test shows that the model satisfies the variance stationarity condition. The density of the unconditional variance $Var(\varepsilon_t)$ can be estimated using $\alpha_0/(1-\alpha_1)$ if required.

7.2 Model validations

Modelling validation for the Brisbane weather station are shown in Figure 2. Panel (a) is a Q-Q plot where empirical quantiles are plotted on the y-axis, and modelled quantiles are on the x-axis. The empirical return periods of the wet season extreme weekly rainfall in Brisbane is presented in Panel (b) with the 95% CI given as dashed lines. We see that the ARMA-GARCH-GH model not only fits the tails of the rainfall distribution but fits the whole distribution very well. Furthermore, the model captures the dependence among rainfall observations. Panel (c) plots the autocorrelation function of the data simulated by the ARMA-GARCH-GH model against the observed autocorrelation function of the rainfall data. We notice that the autocorrelation of two datasets almost align for small lags. Since the model can efficiently represent short-term dependencies in the data, it is suitable for evaluating duration over thresholds. Panel (d) shows that the model produces a good estimate of the distributions of numbers of weeks with rainfall over chosen thresholds, comparing to the empirical distributions. Panel (e) is a plot of duration of very light rainfall, showing the distributions of numbers of consecutive weeks with rainfall less than certain amount. Finally Panel (f) shows durations of zero rainfall. It illustrates the numbers of consecutive weeks with no rain as a bar plot. Overall the ARMA-GARCH-GH model fits the rainfall data well.

8 Comparison to the Generalised Pareto and Weibull Distribution

We compare our censored ARMA-GARCH-GH method to the widely used approach based on a generalised Pareto distribution. Using the method adopted by Li et al. (2005), we choose thresholds $u = 150, 200$ and 250 for the GP distribution. We compare the results of the Pareto model with the ARMA-GARCH-GH model in the top panel of Figure 3, both applied to the Brisbane weather station. It shows that the CI estimates vary when the threshold changes, but the GH model always produces narrower CI in the tails. This indicates that useful information from the centre of the data is being discarded for estimating extreme rainfall return periods by the GP approach. Similar observations can be made when we compare the CIs for return periods produced by the GH and the GW model for the Brisbane weather station. It is clear that the generalised hyperbolic model has narrower CI for rainfall extremes. Furthermore, both the GP and GW methods focus solely on distributional properties of rainfall and rainfall extremes, and they fail to account for serial dependencies in rainfall data. In the ARMA-GARCH-GH model, however, the autocorrelation and heteroscedasticity are addressed by the ARMA and GARCH components so that short-term dependencies in the data can be efficiently captured. Thus it is suitable for evaluating duration over thresholds as well as low rainfall and dry periods. Similar results are observed for rainfall data from weather stations at the Oaks and Pardelup (Appendix B).

9 Discussion

The analysis presented in this paper demonstrates that the ARMA-GARCH-GH model has the potential to outperform conventional methods, such as the GP and the GW distribution for modelling rainfall and rainfall extremes. It appears that useful information from the centre of the data is being discarded for estimating extreme rainfall return periods by the GP approach. Our analysis also

shows that the ARMA-GARCH-GH model not only fits the tails of the distribution, but fits the whole distribution well. The model uses an ARMA and GARCH structure to efficiently represents short-term dependencies in the data. This modelling approach is well suited for a number of applications where there is demand for return periods of extreme rainfall, durations of heavy and light rainfall, as well as dry periods. In its comparison with the GP distribution, we see that the ARMA-GARCH-GH model produces much narrower CIs in the tails. Similarly it also outperformed the GW distribution in its ability to fit the entire rainfall distribution. Furthermore, the ARMA-GARCH-GH model allows us to evaluating duration over and below thresholds partly because it accounts for short-term dependencies in the data, and because it has an effective mechanism for dealing with zero rainfall through censoring.

There are various ways that the methods presented in this paper could be developed further to better fulfil application requirements. There has been a surge of interest in space-time modelling of rainfall and rainfall extremes. Hierarchical models based on well known time series modelling methods such as the dynamical linear models and the AR models are often used in the literature (Bakar and Sahu, 2013; Sahu et al., 2010). The development of a spatio-temporal version of the ARMA-GARCH-GH is currently being undertaken.

Acknowledgements

The authors wish to acknowledge funding for this research from the Digital Productivity and Services Flagship, CSIRO. This work was also financially supported by AusAID. The authors would also like to thank Warren Jin, Steven Crimp and Kristen Williams who provided insightful comments on this paper.

A Proofs of propositions

A.1 Proposition 1

Proof. We want to show that the GH distribution can match the tail decay of the generalised Pareto distribution given by

$$\lim_{x\to\infty} d_{GP}(x) \sim x^{-\left(\frac{1}{\kappa}+1\right)}, \kappa \neq 0. \tag{20}$$

We first fix $\Sigma = 1$ in the GH parameterisation to avoid the identification problem. To get some intuition on the proof, let us consider the reparametrisation

$$\eta = \rho\xi, \;\; \xi = \frac{1}{\sqrt{1+\zeta}}, \text{ where } \rho = \frac{\gamma}{\sqrt{\psi+\gamma^2}}, \;\; \zeta = \sqrt{\chi\psi}.$$

The (η, ξ) parameterisation provides a useful view of the shape of the GH distribution. The parameter restrictions imply $0 < |\eta| < \xi < 1$, so all possible values for η and ξ lie in the interior of a triangle with corners $(-1,1)$, $(1,1)$ and $(0,0)$. This is the hyperbolic *shape triangle*. As it approaches the right hand side of the shape triangle, the distribution becomes more right-skewed with heavier tails. Thus we borrow the ideas of Eberlein and Hammerstein (2004) and approximate

the right hand side boundary of the shape triangle, then we will investigate the properties of this limiting distribution as $x \to \infty$. One obvious way to achieve this is to let $\sqrt{\psi+\gamma^2} \to \infty$ and $\chi \to 0$. We also assume

$$\sqrt{\psi+\gamma^2}-\gamma = \delta > 0, \tag{21}$$

$$\chi\sqrt{\psi+\gamma^2} = \tau > 0. \tag{22}$$

To do this, we need to make use of several asymptotic properties of the Bessel function K_v, we write them down here explicitly:

$$K_v(x) \sim \frac{1}{2}\Gamma(v)\left(\frac{x}{2}\right)^{-v}, \text{ for } v > 0, x \to 0, \tag{23}$$

$$K_v(x) \sim \frac{1}{2}\Gamma(-v)\left(\frac{x}{2}\right)^{v}, \text{ for } v < 0, x \to 0, \tag{24}$$

$$K_v(x) \sim \sqrt{\frac{\pi}{2x}}e^{-x}, \ x \to \infty, \tag{25}$$

$$K_v(x) \sim -\ln(x), \text{ for } v \to 0. \tag{26}$$

Using (25), for $\sqrt{\psi+\gamma^2}$ large enough, we have

$$\frac{K_{\lambda-\frac{1}{2}}\left(\sqrt{\left(\chi+(x-\mu)^2\right)(\psi+\gamma^2)}\right)}{\sqrt{2\pi}\left(\sqrt{\left(\chi+(x-\mu)^2\right)(\psi+\gamma^2)}\right)^{\frac{1}{2}-\lambda}} \sim \frac{\left(\sqrt{\left(\chi+(x-\mu)^2\right)(\psi+\gamma^2)}\right)^{\lambda-1}}{2\exp\left(\sqrt{\left(\chi+(x-\mu)^2\right)(\psi+\gamma^2)}\right)}. \tag{27}$$

Then

$$d_{GH}(x) \sim \frac{\left(\sqrt{\frac{\psi}{\chi}}\right)^{\lambda}\left(\psi+\gamma^2\right)^{\frac{1}{2}-\lambda}\exp\left(\gamma(x-\mu)-\sqrt{\left(\chi+(x-\mu)^2\right)(\psi+\gamma^2)}\right)}{2K_\lambda\left(\sqrt{\chi\psi}\right)\left(\sqrt{\left(\chi+(x-\mu)^2\right)(\psi+\gamma^2)}\right)^{1-\lambda}}$$

$$= \frac{\psi^{\lambda/2}\left(\chi^{-\lambda/2}\left(\psi+\gamma^2\right)^{-\lambda/4}\right)\exp\left(\gamma(x-\mu)-\sqrt{\left(\chi+(x-\mu)^2\right)(\psi+\gamma^2)}\right)}{(\psi+\gamma^2)^{\lambda/4}\,2K_\lambda\left(\sqrt{\chi\psi}\right)\left(\chi+(x-\mu)^2\right)^{(1-\lambda)/2}} \tag{28}$$

$$\sim \frac{\left(\frac{2\delta}{\tau}\right)^{\lambda/2}\exp\left(\gamma(x-\mu)-\sqrt{(\psi+\gamma^2)}\sqrt{\left(\chi+(x-\mu)^2\right)}\right)}{2K_\lambda\left(\sqrt{\chi\psi}\right)\left(\chi+(x-\mu)^2\right)^{(1-\lambda)/2}} \tag{29}$$

$$\sim \frac{\left(\frac{2\delta}{\tau}\right)^{\lambda/2}\exp\left(\gamma(x-\mu)-\sqrt{(\psi+\gamma^2)}\,|x-\mu|\right)}{2K_\lambda\left(\sqrt{2\delta\tau}\right)|x-\mu|^{1-\lambda}}. \tag{30}$$

From (28) to (29) we used, for $\sqrt{\psi+\gamma^2}\to\infty$

$$\frac{\psi^{\lambda/2}}{(\psi+\gamma^2)^{\lambda/4}}=\frac{\left(\left(\sqrt{\psi+\gamma^2}\right)^2-\gamma^2\right)^{\lambda/2}}{\sqrt{\psi+\gamma^2}^{\lambda/2}}=\frac{\left(\left(\sqrt{\psi+\gamma^2}\right)^2-\left(\sqrt{\psi+\gamma^2}-\delta\right)^2\right)^{\lambda/2}}{\sqrt{\psi+\gamma^2}^{\lambda/2}}$$
$$=\left(2\delta-\frac{\delta^2}{\sqrt{\psi+\gamma^2}}\right)^{\lambda/2}\to(2\delta)^{\lambda/2}. \tag{31}$$

From (29) to (30) we used

$$K_\lambda\left(\sqrt{\chi\psi}\right)=K_\lambda\left(\sqrt{\chi}\sqrt{\left(\sqrt{\psi+\gamma^2}\right)^2-\gamma^2}\right)=K_\lambda\left(\sqrt{\chi}\sqrt{\left(\sqrt{\psi+\gamma^2}\right)^2-\left(\sqrt{\psi+\gamma^2}-\delta\right)^2}\right)$$
$$=K_\lambda\left(\sqrt{2\delta\chi\sqrt{\psi+\gamma^2}-\chi\delta^2}\right)\to K_\lambda\left(\sqrt{2\delta\tau}\right),\ \text{for}\,\chi\to 0. \tag{32}$$

Using (21), it follows from (30) that, for $x>\mu$

$$d_{GH}(x)\sim\frac{\left(\frac{2\delta}{\tau}\right)^{\lambda/2}\exp\left(-\delta\left(x-\mu\right)\right)\left(x-\mu\right)^{\lambda-1}}{2K_\lambda\left(\sqrt{\delta\tau}\right)}. \tag{33}$$

To finally obtain the limiting distribution as $\delta\to 0$, we consider the following cases.
Case 1. For $\lambda\in\mathbb{R}_-$. Using (24), (33) can be rewritten as

$$d_{GH}(x)\sim\frac{\left(\frac{2\delta}{\tau}\right)^{\lambda/2}\left(x-\mu\right)^{\lambda-1}\exp\left(-\delta\left(x-\mu\right)\right)}{\Gamma(-\lambda)}\left(\frac{2}{\sqrt{\delta\tau}}\right)^{\lambda}=\frac{\left(\frac{4}{\tau}\right)^{\lambda}\left(x-\mu\right)^{\lambda-1}\exp\left(-\delta\left(x-\mu\right)\right)}{\Gamma(-\lambda)}. \tag{34}$$

Consequently

$$\lim_{\delta\to 0}d_{GH}(x)\sim\frac{\left(\frac{4}{\tau}\right)^{\lambda}\left(x-\mu\right)^{\lambda-1}}{\Gamma(-\lambda)}, \tag{35}$$

which is well-defined for $\lambda\nrightarrow -\infty$. Now, let $x\to\infty$ to get

$$\lim_{x\to\infty}\lim_{\delta\to 0}d_{GH}(x)\sim\frac{\left(\frac{4}{\tau}\right)^{\lambda}x^{\lambda-1}}{\Gamma(-\lambda)}. \tag{36}$$

To match its tail decay with (20), we only require $\lambda=-\frac{1}{\kappa}$ for $\kappa\in\mathbb{R}_+$.
Case 2. For $\lambda=0$. Using (26), we have for $\delta\to 0$

$$K_\lambda\left(\sqrt{\delta\tau}\right)\to-\ln\left(\sqrt{\delta\tau}\right)\to\infty, \tag{37}$$

so $\lim_{\delta\to\infty}d_{GH}(x)\to 0$, and there is no limiting probability distribution. We should exclude this case.

Case 3. For $\lambda \in \mathbb{R}_+$. Using (23), we have for $\delta \to 0$

$$d_{GH}(x) \sim \frac{\left(\frac{2\delta}{\tau}\right)^{\lambda/2} x^{\lambda-1}\exp(-\delta x)}{\Gamma(\lambda)}\left(\frac{\sqrt{\delta\tau}}{2}\right)^{\lambda} = \frac{\left(\frac{\delta}{\sqrt{2}}\right)^{\lambda} x^{\lambda-1}\exp(-\delta x)}{\Gamma(\lambda)} \to 0. \tag{38}$$

Hence there is no limiting distribution. However, it is well-known that generalised Pareto distirbutions converge to the exponential function as $\kappa \to 0$, i.e.

$$\lim_{x\to\infty}\lim_{\kappa\to 0} d_{GP(\kappa,\mu,\sigma)}(x) \sim \frac{1}{\sigma}\exp\left(-\frac{x}{\sigma}\right). \tag{39}$$

In this case, we do not need the limiting distribution as $\delta \to 0$. The exponential decay of (39) can be easily achieved by a hyperbolic distribution (i.e. $\lambda = 1$), we have

$$\lim_{x\to\infty} d_{GH}(x) \sim \frac{\left(\frac{2\delta}{\tau}\right)^{1/2}\exp(-\delta x)}{2K_\lambda\left(\sqrt{\delta\tau}\right)}. \tag{40}$$

□

A.2 Proposition 2

Proof. We will make use of mean and variance expressions of the generalised inverse Gaussian distirbution, whose density function is given by

$$d_{GIG(\lambda,\chi,\psi)}(x) = \left(\frac{\sqrt{\psi}}{\sqrt{\chi}}\right)^{\lambda}\frac{1}{2K_\lambda\left(\sqrt{\chi\psi}\right)}x^{\lambda-1}\exp\left(-\frac{1}{2}\left(\frac{\chi}{x}+\psi x\right)\right),\ x>0, \tag{41}$$

where $\chi, \psi > 0$ and $\lambda \in \mathbb{R}$. It has the moment generating function

$$E\left(X^k\right) = \left(\frac{\sqrt{\chi}}{\sqrt{\psi}}\right)^{k}\frac{K_{\lambda+k}\left(\sqrt{\chi\psi}\right)}{K_\lambda\left(\sqrt{\chi\psi}\right)}. \tag{42}$$

Now for $X \sim GH(\lambda,\chi,\psi,\mu,\Sigma,\gamma)$, it follows from (6) and (42) that if we impose $Var(X) = 1$, we then get

$$\frac{\gamma^2}{\psi}\left(\frac{K_{\lambda+2}\left(\sqrt{\psi}\right)}{K_\lambda\left(\sqrt{\psi}\right)} - \frac{K^2_{\lambda+1}\left(\sqrt{\psi}\right)}{K^2_\lambda\left(\sqrt{\psi}\right)}\right) + \frac{\Sigma}{\sqrt{\psi}}\frac{K_{\lambda+1}\left(\sqrt{\psi}\right)}{K_\lambda\left(\sqrt{\psi}\right)} = 1, \tag{43}$$

Solving this simple linear equation for Σ, we have

$$\Sigma = \frac{\sqrt{\psi}}{M_{\lambda+1}\left(\sqrt{\psi}\right)} - \frac{\gamma^2}{\psi}\left(M_{\lambda+2}\left(\sqrt{\psi}\right) - M_{\lambda+1}\left(\sqrt{\psi}\right)\right), \tag{44}$$

where $M_{\lambda+1}\left(\sqrt{\psi}\right) = K_{\lambda+1}\left(\sqrt{\psi}\right)/K_\lambda\left(\sqrt{\psi}\right)$. However we notice that the expression for Σ is not always postive. To overcome this problem, we replace the parameter γ in (4) with the parameter

$\beta = \gamma\Sigma^{-1}$ to obtain the new parameterisation $X \sim GH(\lambda, \chi, \psi, \mu, \Sigma, \beta)$. Then (43) becomes

$$\Sigma^2 \frac{\beta^2}{\psi}\left(\frac{K_{\lambda+2}\left(\sqrt{\psi}\right)}{K_\lambda\left(\sqrt{\psi}\right)} - \frac{K^2_{\lambda+1}\left(\sqrt{\psi}\right)}{K^2_\lambda\left(\sqrt{\psi}\right)}\right) + \frac{\Sigma}{\sqrt{\psi}}\frac{K_{\lambda+1}\left(\sqrt{\psi}\right)}{K_\lambda\left(\sqrt{\psi}\right)} - 1 = 0, \tag{45}$$

whose solutions are

$$\Sigma = \frac{-1 \pm \sqrt{1 + 4\left(N_{\lambda+2}\left(\sqrt{\psi}\right) - 1\right)\beta^2}}{2M_{\lambda+1}\left(\sqrt{\psi}\right)\left(N_{\lambda+2}\left(\sqrt{\psi}\right) - 1\right)\beta^2/\sqrt{\psi}}, \tag{46}$$

where $N_{\lambda+2}\left(\sqrt{\psi}\right) = K_{\lambda+2}\left(\sqrt{\psi}\right)K_\lambda\left(\sqrt{\psi}\right)/K^2_{\lambda+1}\left(\sqrt{\psi}\right)$.

We know that $M_\nu(x) > 0$ for all ν and $x > 0$, and Ismail and Muldoon (1978) showed that $N_\nu(x) > 1$ for $\lambda \in \mathbb{R}$ and $x > 0$, so (45) always has a postive and a negative solutions. Since $\Sigma > 0$, we need to choose

$$\Sigma = \frac{-1 + \sqrt{1 + 4\left(N_{\lambda+2}\left(\sqrt{\psi}\right) - 1\right)\beta^2}}{2M_{\lambda+1}\left(\sqrt{\psi}\right)\left(N_{\lambda+2}\left(\sqrt{\psi}\right) - 1\right)\beta^2/\sqrt{\psi}} \text{ and } \mu = -\frac{\beta\Sigma M_{\lambda+1}\left(\sqrt{\psi}\right)}{\sqrt{\psi}}.$$

□

B Modelling results for the Oaks and Pardelup

References

Ashcroft, M., K. French, and L. Chisholm (2011). An evaluation of environmental factors affecting species distributions. *Ecological Modelling 222*, 524–531.

Ashcroft, M. and J. Gollan (2012). Fine-resolution (25 m) topoclimatic grids of near-surface (5 cm) extreme temperatures and humidities across various habitats in a large (200 300 km) and diverse region. *International Journal of Climatology 32*, 2134–2148.

Bakar, K. S. and P. Kokic (2014). A spatially varying approach for precipitation modelling in south-western australia.

Bakar, K. S. and S. K. Sahu (2013). sptimer: Spatio-temporal bayesian modelling using r.

Barndorff-Nielsen, O. (1977). Exponentially decreasing distributions for the logarithm of particle size. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences 353*(1674), 401–419.

Barndorff-Nielsen, O. E. and N. Shephard (2001). *Normal modified stable processes*. MaPhySto, Department of Mathematical Sciences, University of Aarhus.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics 31*(3), 307–327.

Breymann, W. and D. Lüthi (2013). ghyp: A package on generalized hyperbolic distributions.

Burr, I. W. (1942). Cumulative frequency functions. *The Annals of Mathematical Statistics 13*(2), 215–232.

Champernowne, D. G. (1952). The graduation of income distributions. *Econometrica: Journal of the Econometric Society*, 591–615.

Coles, S., J. Bawa, L. Trenner, and P. Dorazio (2001). *An introduction to statistical modeling of extreme values*, Volume 208. Springer.

Coles, S., L. R. Pericchi, and S. Sisson (2003). A fully probabilistic approach to extreme rainfall modeling. *Journal of Hydrology 273*(1), 35–50.

Davison, A. C. and R. L. Smith (1990). Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B (Methodological)*, 393–442.

Drosdowsky, W. and L. E. Chambers (2001). Near-global sea surface temperature anomalies as predictors of australian seasonal rainfall. *Journal of Climate 14*(7), 1677–1687.

DuMouchel, W. H. (1983). Estimating the stable index α in order to measure tail thickness: a critique. *The Annals of Statistics*, 1019–1031.

Eberlein, E. and E. A. v. Hammerstein (2004). Generalized hyperbolic and inverse gaussian distributions: limiting cases and approximation of processes. In *Seminar on Stochastic Analysis, Random Fields and Applications IV*, pp. 221–264. Springer.

Finkenstadt, B. and H. Rootzén (2004). *Extreme values in finance, telecommunications, and the environment*. CRC Press.

Fowler, H. J., S. Blenkinsop, and C. Tebaldi (2007). Linking climate change modelling to impacts studies: recent advances in downscaling techniques for hydrological modelling. *International Journal of Climatology 27*, 1547–1578.

Gilli, M. et al. (2006). An application of extreme value theory for measuring financial risk. *Computational Economics 27*(2-3), 207–228.

Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika 57*(1), 97–109.

Hosking, J. R. and J. R. Wallis (1987). Parameter and quantile estimation for the generalized pareto distribution. *Technometrics 29*(3), 339–349.

Hutchinson, M. (1995). Stochastic space-time weather models from ground-based data. *Agricultural and Forest Meteorology 73*, 237–264.

Ismail, M. E. and M. E. Muldoon (1978). Monotonicity of the zeros of a cross-product of bessel functions. *SIAM Journal on Mathematical Analysis 9*(4), 759–767.

Joe, H. (1987). Estimation of quantiles of the maximum of n observations. *Biometrika 74*(2), 347–354.

Keating, B. A., P. S. Carberry, G. L. Hammer, M. E. Probert, M. J. Robertson, D. Holzworth, N. I. Huth, J. N. Hargreaves, H. Meinke, Z. Hochman, and et al. (2003). An overview of APSIM, a model designed for farming systems simulation. *European Journal of Agronomy 18*, 267288.

King, A. D., L. V. Alexander, and M. G. Donat (2013). Asymmetry in the response of eastern australia extreme rainfall to low-frequency pacific variability. *Geophysical Research Letters 40*(10), 2271–2277.

Kleiber, C. and S. Kotz (2003). *Statistical size distributions in economics and actuarial sciences*, Volume 470. John Wiley & Sons.

Kokic, P., H. Jin, and S. Crimp (2013). Improved point scale climate projections using a block bootstrap simulation and quantile matching method. *Climate Dynamics 41*, 853–866.

K.S.Bakar, a. P. and H. Jin (2014). A spatio-dynamic model for assessing frost risk in south-east australia.

Li, Y., W. Cai, and E. Campbell (2005). Statistical modeling of extreme rainfall in southwest western australia. *Journal of Climate 18*(6).

Maddala, G. S. (1986). *Limited-dependent and qualitative variables in econometrics*. Number 3. Cambridge university press.

Mencia, J. and E. Sentana (2004). Estimation and testing of dynamic models with generalised hyperbolic innovations.

Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics 21*(6), 1087–1092.

Mudholkar, G. S. and A. D. Huston (1996). The exponentiated Weibull family: some properties and a flood data application. *Communication in Statistics - Theory and Methods 25*, 3059–3083.

Pickands III, J. (1971). The two-dimensional poisson process and extremal processes. *Journal of Applied Probability*, 745–756.

Sahu, S. K., A. E. Gelfand, and D. M. Holland (2010). Fusing point and areal level space–time data with application to wet deposition. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 59*, 77–103.

Schepen, A., Q. Wang, and D. Robertson (2012). Evidence for using lagged climate indices to forecast australian seasonal rainfall. *Journal of Climate 25*(4), 1230–1246.

Shao, Q., Q. Wang, and L. Zhang (2013). A stochastic weather generation method for temporal precipitation simulation. In *20th International Congress on Modelling and Simulation, Adelaide, Australia, 16 December, 2013*, pp. 2681–2687.

Shao, Q., H. Wong, J. Xia, and I. W.-C. (2004). Models for extremes using the extended three parameter Burr XII system with application to flood frequency analysis. *Hydrological Sciences Journal 49*, 685–702.

Smith, R. L. (1984). Threshold methods for sample extremes. In *Statistical extremes and applications*, pp. 621–638. Springer.

Smith, R. L. (1987). Estimating tails of probability distributions. *The annals of Statistics*, 1174–1207.

Stone, R. C., G. L. Hammer, and T. Marcussen (1996). Prediction of global rainfall probabilities using phases of the southern oscillation index.

Tetzlaff, D., S. Uhlenbrook, and P. Molnar (2005). Significance of spatial variability in precipitation for process-oriented modelling: results from two nested catchments using radar and ground station data. *Hydrology & Earth System Sciences 9*.

Van Montfort, M. and J. Witter (1986). The generalized pareto distribution applied to rainfall depths. *Hydrological Sciences Journal 31*(2), 151–162.

Wang, Q. J., D. L. Shrestha, D. E. Robertson, and P. Pokhrel (2012). A log-sinh transformation for data normalization and variance stabilization. *Water Resources Research 48*, 1–7.

Whittle, P. (1951). *Hypothesis Testing in Time Series Analysis*. Ph. D. thesis, University of Uppsala.

Ye, W. and Y. Li (2011). A method of applying daily gcm outputs in assessing climate change impact on multiple day extreme precipitation for brisbane river catchment. In *19th international Congress on modelling and simulation*, pp. 12–16.

Yiou, P., K. Goubanova, Z. Li, and M. Nogaj (2008). Weather regime dependence of extreme value statistics for summer temperature and precipitation. *Nonlinear Processes in Geophysics 15*(3).

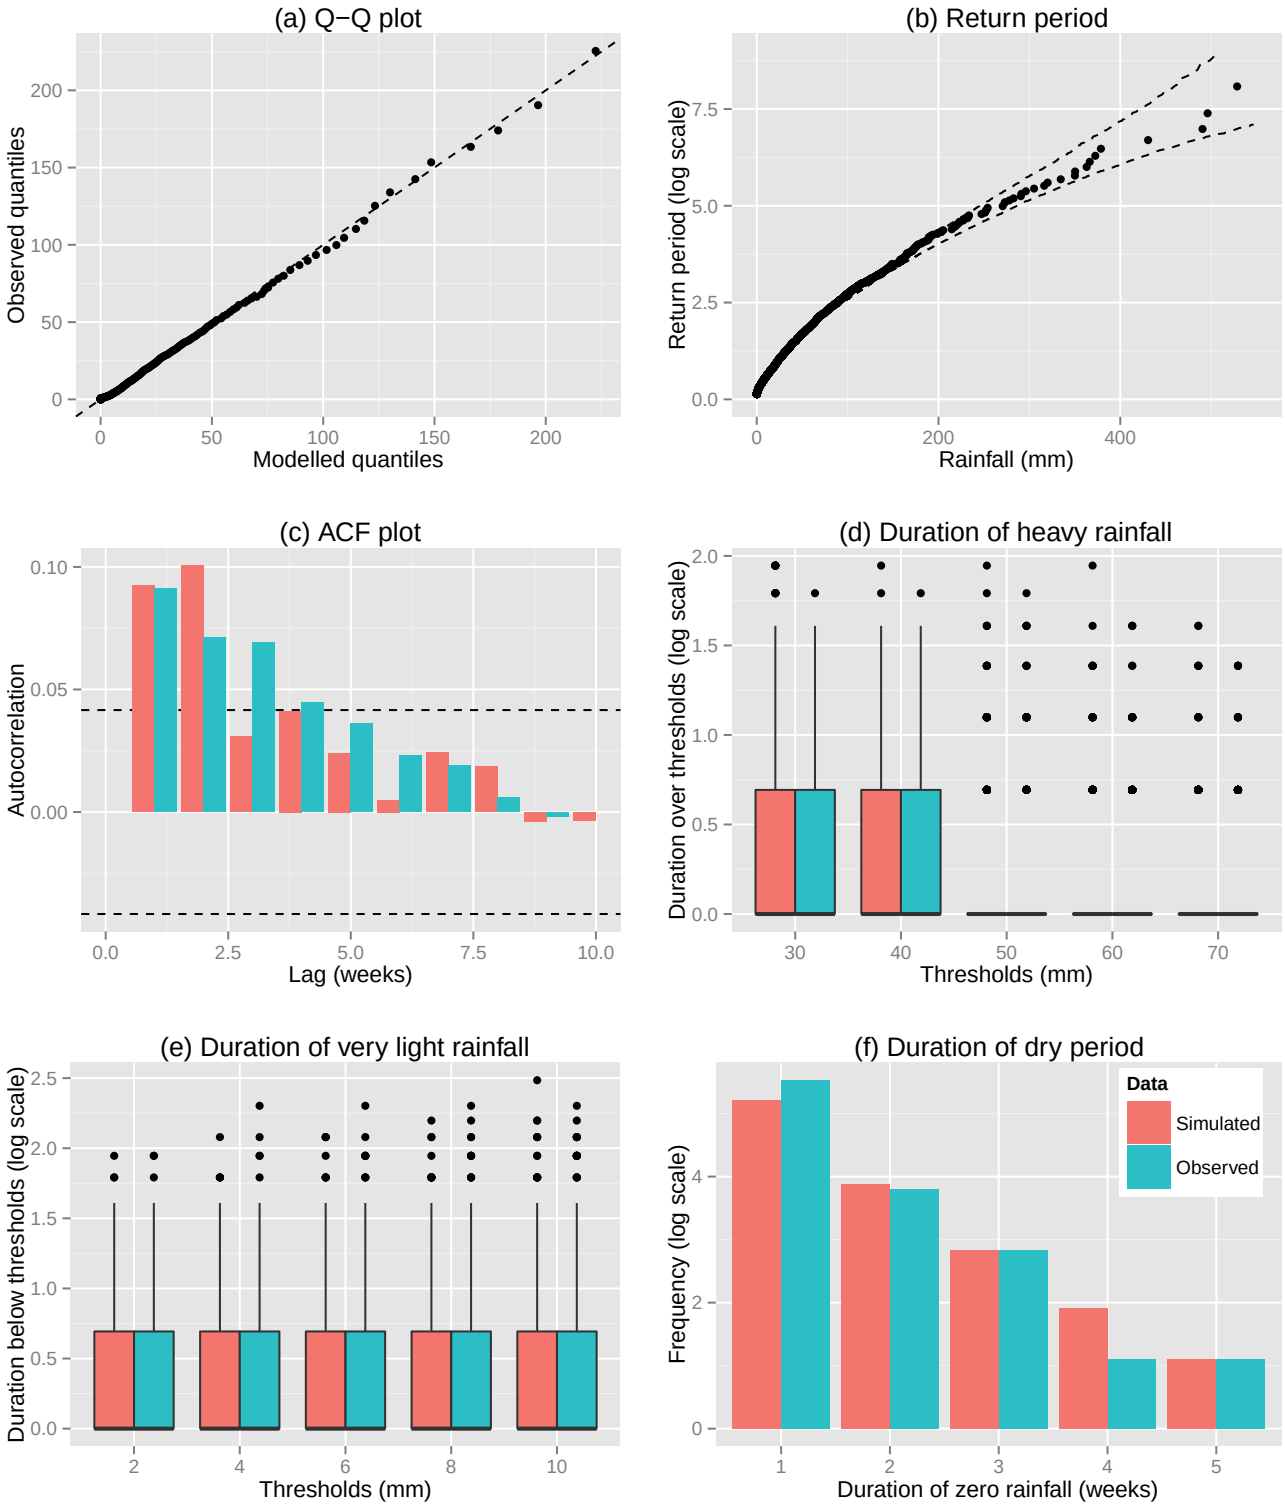

Figure 2: Modelling results for the weather station Brisbane. Panel (a) is a Q-Q plot where empirical quantiles are plotted on the y-axis, and modelled quantiles are on the x-axis. Panel (b) shows return periods of the wet season weekly rainfall. The dots are empirical return periods, and the 95% CI given as dashed lines. Panel (c) shows the autocorrelation function of the data simulated by the model comparing to the observed autocorrelation function of the rainfall data. Panel (d) shows the distributions of numbers of weeks with simulated rainfall over chosen thresholds, comparing to the empirical distributions. Panel (e) is a plot of duration of very light rainfall, showing the distributions of numbers of consecutive weeks with rainfall less than certain amount, also comparing to the empirical distribution. Finally Panel (f) shows durations of zero rainfall for both observed and simulated data. All based on 1,000 simulations.

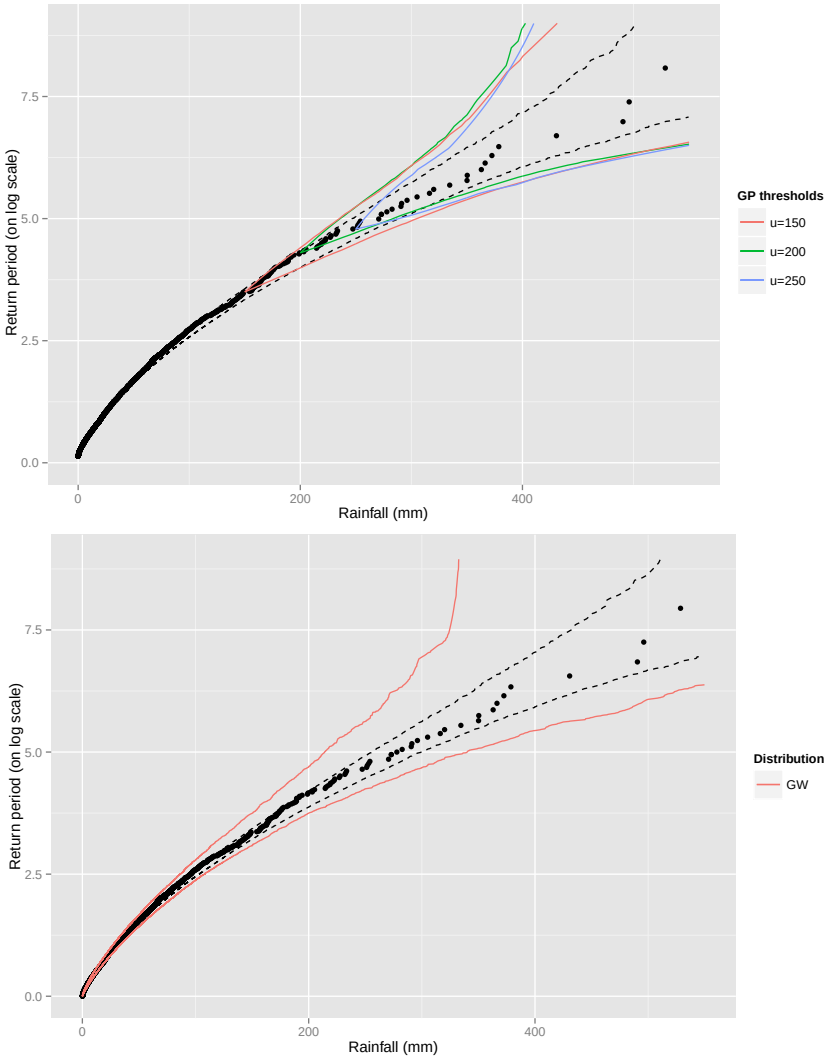

Figure 3: Return periods of the extreme winter weekly rainfall at Brisbane station. The broken curves represent the 95% CI estimated by the ARMA-GARCH-GH model. In the left panel, the solid curves are the 95% CIs produced by the Pareto model with different thresholds. In the right panel, the solid curves are the 95% CIs produced by the generalised Weibull distribution.

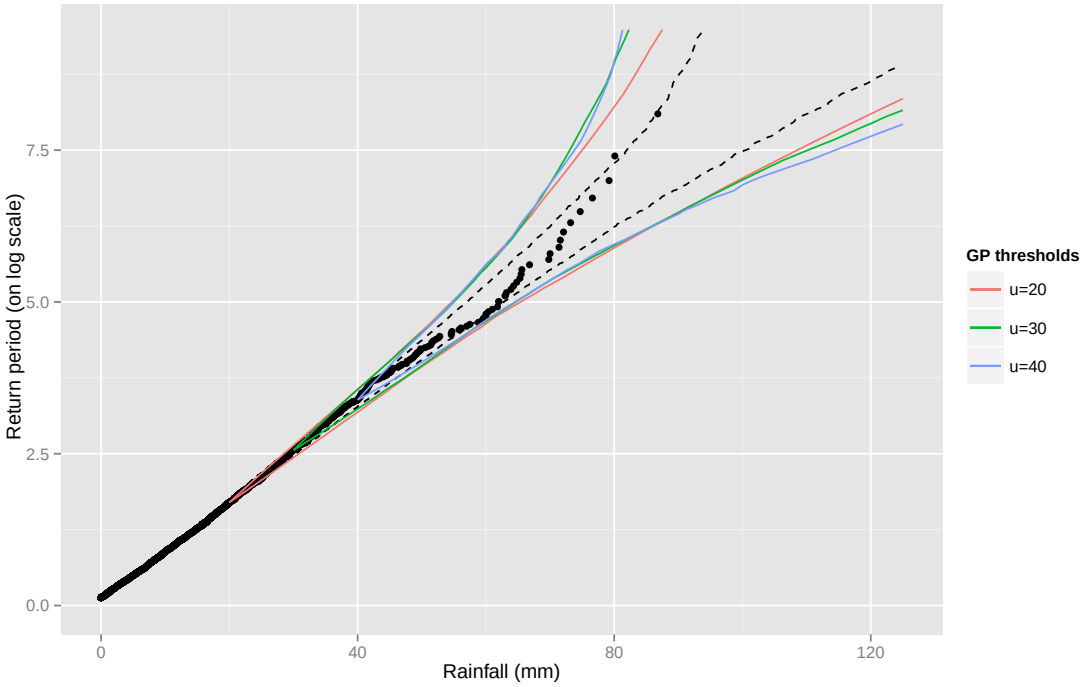

Figure 4: Return periods of winter extreme weekly rainfall at the Oaks weather station, based on 1,000 simulations. The broken curves represent the 95% CI estimated by the ARMA-GARCH-GH model, and the solid curves are the 95% CIs produced by the Pareto model with different thresholds.

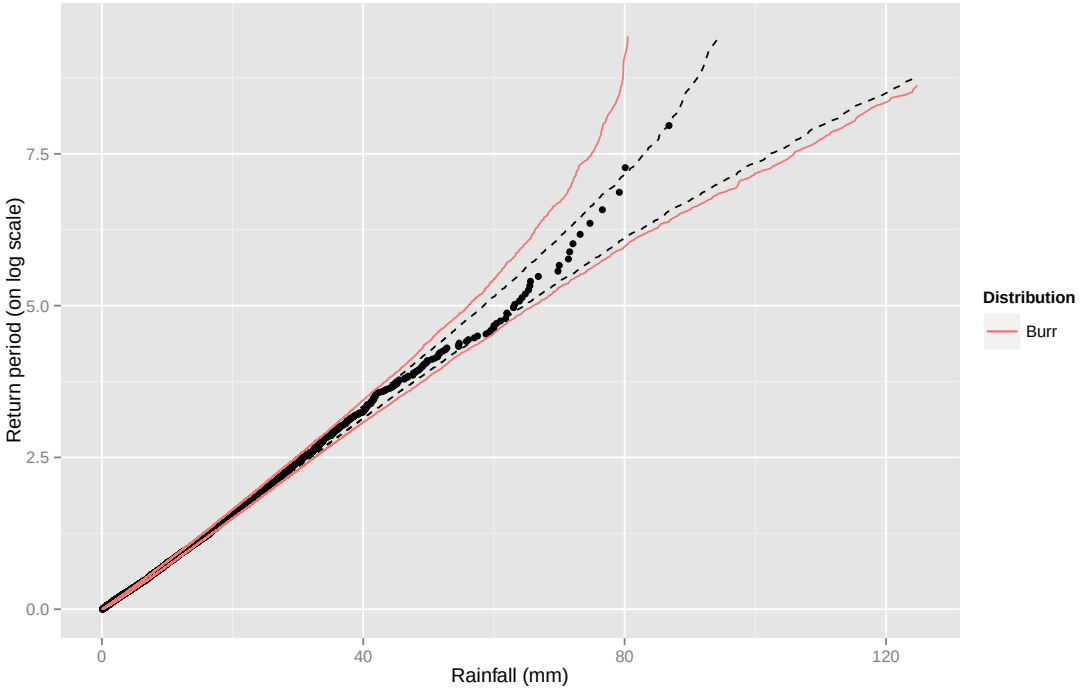

Figure 5: Return periods of winter extreme weekly rainfall at the Oaks weather station, based on 1,000 simulations. The broken curves represent the 95% CI estimated by the ARMA-GARCH-GH model, and the solid curves are the 95% CIs produced by the Burr distribution.

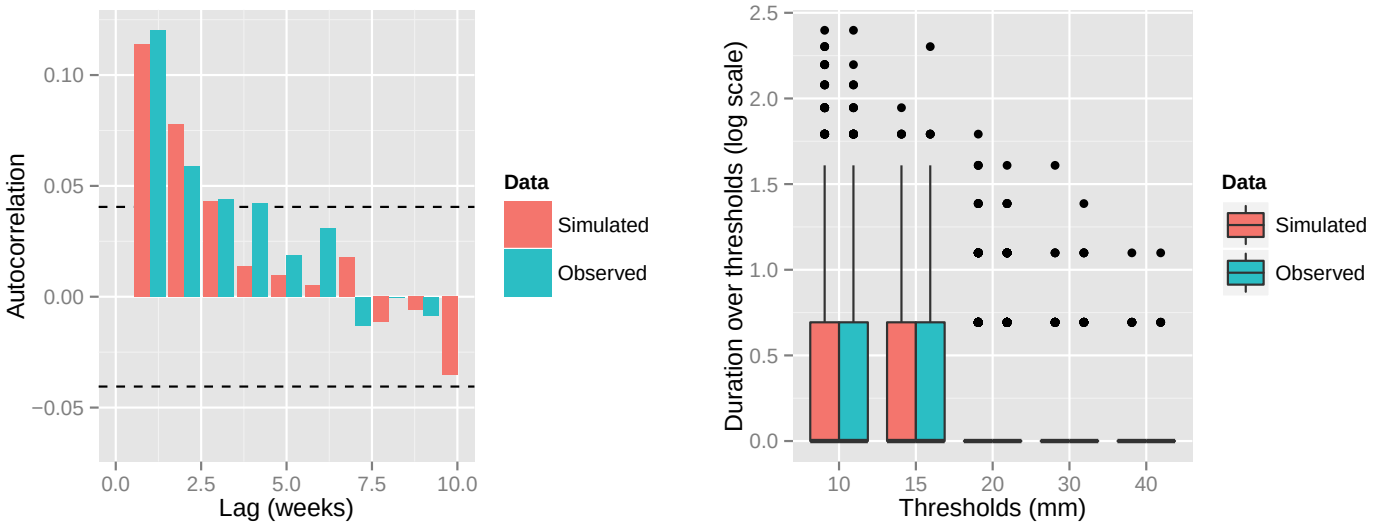

Figure 6: The left panel shows the comparison of autocorrelation in the Oaks data and the data simulated by the ARMA-GARCH-GH model; the right panel is a boxplot (on log scale) of numbers of consecutive weeks with total rainfall over thresholds. Both are based on 1,000 simulations.

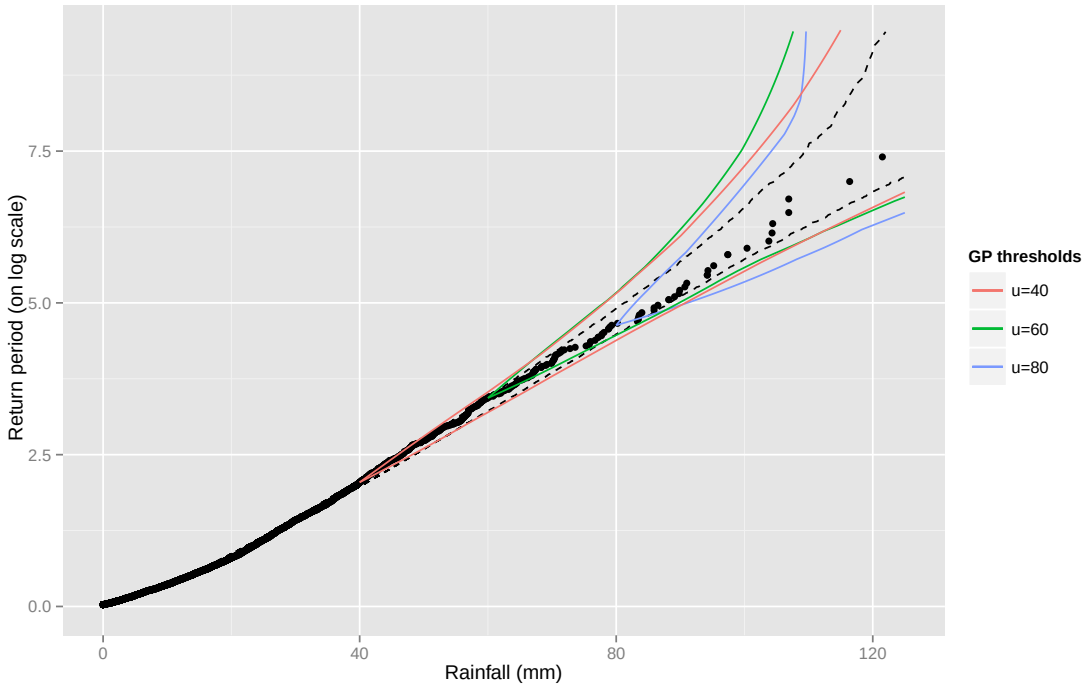

Figure 7: Return periods of winter extreme weekly rainfall at Pardelup weather station, based on 1,000 simulations. The broken curves represent the 95% CI estimated by the ARMA-GARCH-GH model, and the solid curves are the 95% CIs produced by the Pareto model with different thresholds.

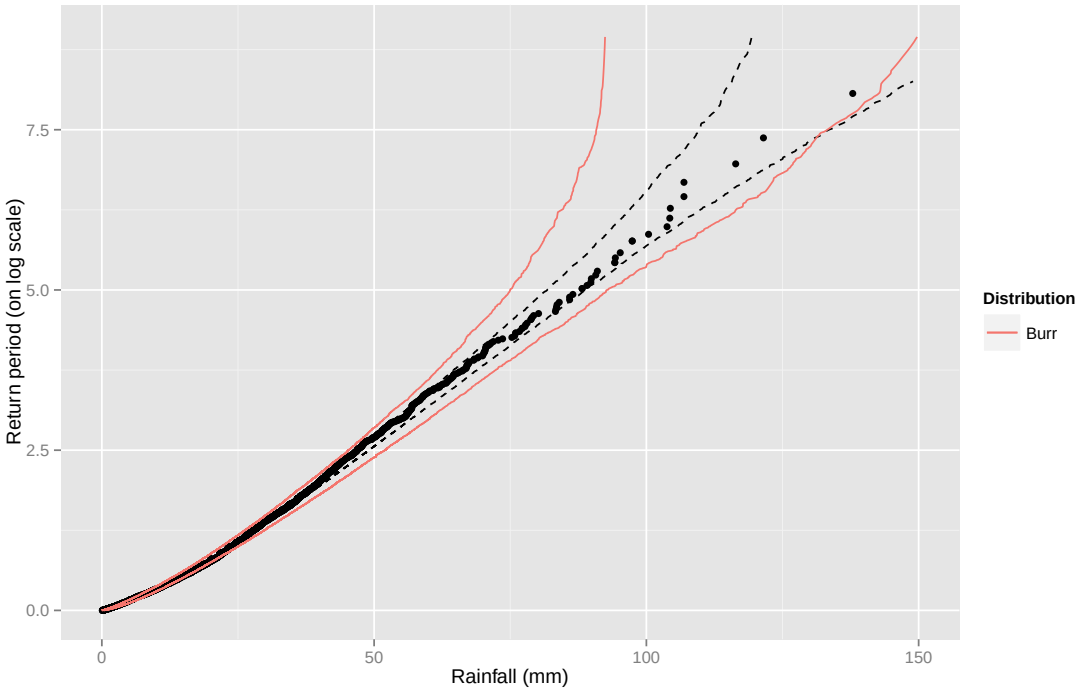

Figure 8: Return periods of winter extreme weekly rainfall at Pardelup weather station, based on 1,000 simulations. The broken curves represent the 95% CI estimated by the ARMA-GARCH-GH model, and the solid curves are the 95% CIs produced by the Burr distribution.

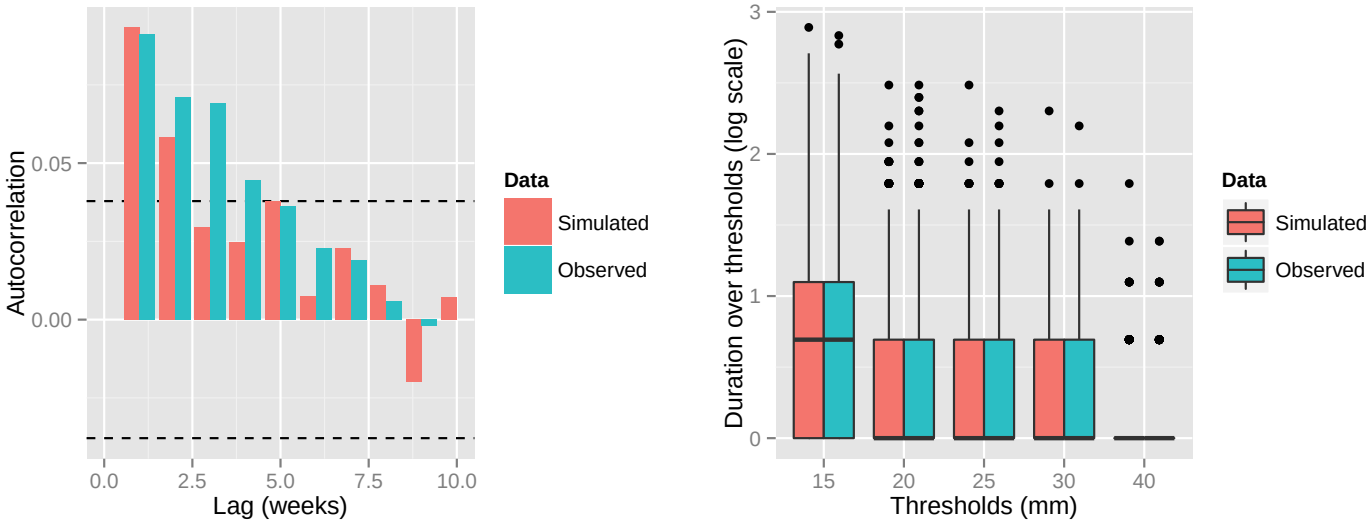

Figure 9: The left panel shows the comparison of autocorrelation in Pardelup data and the data simulated by the ARMA-GARCH-GH model; the right panel is a boxplot (on log scale) of numbers of consecutive weeks with total rainfall over thresholds. Both are based on 1,000 simulations.